

Topological language for RNA

Fenix W.D. Huang^a, Christian M. Reidys^{a,*}

^a*Biocomplexity Institute of Virginia Tech, Virginia Tech, United States of America*

Abstract

In this paper we introduce a novel, context-free grammar, *RNAFeatures**, capable of generating any RNA structure including pseudoknot structures (pk-structure). We represent pk-structures as orientable fatgraphs, which naturally leads to a filtration by their topological genus. Within this framework, RNA secondary structures correspond to pk-structures of genus zero. *RNAFeatures** acts on formal, arc-labeled RNA secondary structures, called λ -structures. λ -structures correspond one-to-one to pk-structures together with some additional information. This information consists of the specific rearrangement of the backbone, by which a pk-structure can be made cross-free. *RNAFeatures** is an extension of the grammar for secondary structures and employs an enhancement by labelings of the symbols as well as the production rules. We discuss how to use *RNAFeatures** to obtain a stochastic context-free grammar for pk-structures, using data of RNA sequences and structures. The induced grammar facilitates fast Boltzmann sampling and statistical analysis. As a first application, we present an $O(n \log(n))$ runtime algorithm which samples pk-structures based on ninety tRNA sequences and structures from the Nucleic Acid Database (NDB).

Keywords:

RNA pseudoknot structure, context-free grammar, topological RNA structure, fatgraph, stochastic context-free grammar

*Corresponding author

Email addresses: `fenixprotoss@gmail.com` (Fenix W.D. Huang),
`duck@santafe.edu` (Christian M. Reidys)

1. Introduction

RNA secondary structure provides a coarse grained model to study important features of the real molecular conformation, embedded in three-space [1]. The key feature of these structures is that they can be inductively constructed [2]. The recursion forms the basis for more than three decades of research resulting in what can be called the dynamic programming (DP) paradigm. DP methodology facilitates *ab initio* folding [3, 4, 5], partition function [6], Boltzmann sampling [7] and context-free grammars (CFGs) [8]. Minimum free energy (mfe) secondary structures can be folded in $O(n^2)$ space and $O(N^3)$ time complexity [3] and their partition function is computed in $O(n^3)$ time [6]. Stochastic folding is based on their CFG, allowing to factor in specific properties of biological structures. Thermodynamic and stochastic secondary structures have been statistically analyzed in [9, 10, 5]. In [11] the secondary structure paradigm has been extended to include gap-structures. Gap-structures include certain, but not all pseudoknot structures (pk-structures). They can be viewed as the most general class to which the DP-paradigm directly applies and their underlying grammar is multiple context-free. The approach allowed folding of gap-structures in $O(n^6)$ time and $O(n^4)$ space complexity and sampling in $O(n^3)$ time.

In this paper we approach the problem of cross-serial interactions of pk-structures from first principles. We employ an abstraction, that is well known to structural biology [12]: namely, a nucleotide is not a dimensionless object but a triangle and base pairs can be abstracted as untwisted or twisted ribbons connecting two such triangles, see Fig. 1. This observation introduces topology as a natural language for biological structures and leads to the fatgraph model [13]. Our main result is to map a novel recursion on unicellular maps [14] into the classic recursion derived in [2]. I.e. we translate a recursive procedure on fatgraphs into an enhanced version of the recursion of secondary structures. This results in the first, unambiguous, context-free grammar for pk-structures, *RNAFeatures**. We present as a first application an $O(n \log(n))$ Boltzmann sampler and stochastic context-free grammar.

In the following we discuss the fatgraph modeling Ansatz in detail. We then provide some background on computational linguistics in the context of RNA structures, enabling us to discuss what we mean by topological language of RNA.

The biophysics of the canonical base pairs **A-U**, **G-C** and **G-U** implies, that the fatgraphs modeling RNA structures are orientable: RNA purine and

pyrimidine bases are modeled as three edges for hydrogen bonding interactions: the Watson-Crick edge, the Hoogsteen edge, and the Sugar edge¹ [12]. A given edge of one base can potentially interact in a plane with any one of the three edges of a second base, and can do so in either cis- or trans- orientation of the glycosidic bonds. This gives rise to 12 basic geometric types with at least two H bonds connecting the bases, see Fig. 1. Accordingly, cis-base pairs, including the canonical base pairs belong to the *cis-Watson-Crick/Watson-Crick* geometry.

Canonical base pairs induce exclusively untwisted ribbons, modeled via orientable *fatgraphs*, i.e. graphs enhanced by replacing vertices by discs and edges by ribbons. Gluing the sides of these ribbons produces the topological quotient space [15], which is a connected sum of tori. As topological genus, g , and orientability are the two determinants of closed surfaces, RNA structures having canonical base pairs are filtered by genus alone. RNA secondary structures [16], are exactly structures of genus zero. Structures of genus $g > 0$ are those exhibiting cross-serial interactions. In a sense, topological genus measures the complexity of a biomolecule, far more subtle than the number of crossing arcs. One may think of pk-structures as *cross-free* drawings on a closed surface of genus $g \geq 0$, instead of a drawing with crossing arcs in the plane², see Fig. 2.

The term *language of RNA* has been used in [17], where the authors study the aforementioned gap-structures [11]. Methods from computational linguistics have been applied to problems in DNA and RNA sequence analysis for decades: early work involved using regular grammars and hidden Markov models (HMMs) in order to model biological sequences. These have been used in sequence analysis [18], such as identifying CpG islands [19], gene prediction [20], pairwise and multiple sequence alignment [19, 21], modeling DNA sequencing errors [22], protein secondary structure prediction, and RNA structural alignment [23]. Regular grammars and HMMs are not well-equipped to model problems in RNA folding because they cannot describe the long-distance correlations of base pairs. Instead, a larger class called *context-free* grammars (CFG) are used. Assigning probabilities to the production rules of CFG produces a stochastic context-free grammar (SCFG),

¹which includes the 29-OH and which has also been referred to as the Shallowgroove edge

²A cross-free drawing in the plane corresponds to an RNA secondary structure.

which, in the context of HMM, allows each state to generate not a single, but any (fixed) number of immediate successors.

By a *topological language of RNA* we mean an unambiguous, CFG that allows us to recursively construct *any* pk-structure using recursions that are induced by cell-surgery on fatgraphs. The latter shall be translated into an enhancement of the classical recursion [2]. This allows us to connect to the existing paradigm by presenting a CFG based on arc-removals and decomposition as production rules³, where the cross-serial interactions are encapsulated in the labels of the arcs and nonterminal symbols, respectively.

Finally we discuss the above mentioned recursion of secondary structures and provide some background on topological RNA structures.

Waterman *et al.* [24, 2, 25, 26] studied the combinatorics and folding of RNA secondary structures. Their diagrams are labeled graphs over the vertex set $[n] = \{1, \dots, n\}$, presented by drawing the vertices on a horizontal line and noncrossing arcs in the upper half-plane. Vertices and arcs correspond to the nucleotides **A**, **G**, **U** and **C** and Watson-Crick (**A-U**, **G-C**) and wobble (**U-G**) base pairs, respectively. The noncrossing arcs of RNA secondary structures allow for a recursive build. Let $S_2(n)$ denotes the number of RNA secondary structures over n nucleotides then we have [24]: $S_2(n) = S_2(n-1) + \sum_{j=0}^{n-3} S_2(n-2-j)S_2(j)$, where $S_2(n) = 1$ for $0 \leq n \leq 2$. Accordingly, RNA secondary structures satisfy a constructive recursion. This relation constitutes the basis for the DP-recursions used for the polynomial time folding [16, 4] and has profound algorithmic implications. The DP-framework has ever since strongly influenced the field of RNA folding [24, 4, 27, 5].

CFGs of secondary structures [4, 5, 28] are utilizing the recursive nature of secondary structures [24]. The grammar has two production rules, displayed in Fig. 3. Its terminal symbols denote a pair of vertices forming an arc as well as an unpaired vertex and the nonterminal symbols are an arbitrary secondary structure and an irreducible structure covered by an external arc. The grammar is unambiguous, i.e. each structure has a unique decomposition path.

Cross-serial interactions or pseudoknots have long been known as important structural elements [29, 30] in RNA. Depicting a contact structure as a diagram, a structure contains a pseudoknot if and only if these arcs cross. Cross-serial interactions are functionally important in tRNAs, RNaseP [31],

³as for CFG of secondary structures

telomerase RNA [32], and ribosomal RNAs [33]. Pseudoknots in plant virus RNAs mimic tRNA structures, and *in vitro* selection experiments have produced pseudoknotted RNA families that bind to the HIV-1 reverse transcriptase [34].

Topological RNA structures have been introduced in [35, 36] and the classification and expansion of pk-structures in terms of the topological genus of an associated fatgraph has been studied by means of matrix theory in [37, 38]. The computation of the genus of a fatgraph is classic [39] and was first applied to RNA structures by [37] and [38]. [40] studies topological RNA structures of higher genus and connects them with Riemann’s Moduli space [36]. In [41] a polynomial time, loop-based folding algorithm of topological RNA structures was given.

The paper is organized as follows: in Section 2 we discuss the idea behind the novel recursions. Thereafter we establish the correspondence between pk-structures and blueprints with λ -structures in Section 3. Then we discuss *RNAFeatures** and finally detail the SCFG in Section 4.

2. Fatgraphs and blueprints.

The contacts of nucleotides within an RNA structure can be represented as a diagram whose arcs can be linearly ordered via their left end-points, i.e. $(i, j) \prec (r, s)$ if and only if $i < r$. Suppose $(i, j) \prec (r, s)$, then (i, j) and (r, s) cross iff $i < r < j < s$ holds.

In the following we shall assume that any diagram contains the arc $(0, n+1)$ (rainbow), which is not accounted for as an edge of the diagram. Rainbows mark the start and endpoint of the 5’-3’ oriented backbone, thus allowing to collapse the backbone into a disc without losing any information.

The passage from diagrams to fatgraphs [42, 13] is obtained by “thickening” the edges into (untwisted) bands or ribbons. Furthermore, each vertex is inflated into a disc as shown in Fig. 4. This inflation reflects the fact that canonical base pairs fix a plane, see Fig. 1. A fatgraph, \mathbb{G} , can thus be viewed as a “drawing” on an orientable surface $X_{\mathbb{G}}$, which is obtained by identifying the sides of the ribbons. \mathbb{G} is a 2-dimensional cell complex over its geometric realization, $X_{\mathbb{G}}$.

The Euler characteristic and topological genus of the surface \mathbb{D} are given by $\chi(\mathbb{D}) = v - e + r$ and $g(\mathbb{D}) = 1 - \frac{1}{2}\chi(\mathbb{D})$, where v , e and r denote the number of discs, ribbons and boundary components in \mathbb{D} [15]. The equivalence

of simplicial and singular homology [43] implies that these combinatorial invariants are topological invariants. This means the genus of the surface $X_{\mathbb{D}}$ provides a filtration of fatgraphs.

A fatgraph can be presented by a pair of permutations [13]. Let $H = [2n]$ denote the set of half-edges, and let σ , α and γ be three permutations over H , where each cycle in σ , α and γ presents a vertex, an edge and a boundary component, respectively. A vertex of degree k is considered as a cycle $v = (i, \sigma(i), \dots, \sigma^{k-1}(i))$. In particular, α is a fixed-point free involution since an edge consist of two half-edges and we have $\gamma = \alpha \circ \sigma$.

Removing the unpaired vertices from the diagram and collapsing the backbone into a disc does not change the Euler characteristic. Therefore, the relation between genus and number of boundary components is solely determined by the number of arcs in the upper half-plane: $2 - 2g - r = 1 - n$, where n is number of arcs and r the number of boundary components. The latter can be computed easily and allows us to obtain the genus of the diagram.

Mapping (H, σ, α) into $\pi((H, \sigma, \alpha)) = (H, \alpha \circ \sigma, \alpha)$, is a bijection, see Fig. 5. π is called the *Poincaré dual* and interchanges boundary components with vertices, preserving topological genus. By construction we shall deal only with fatgraphs over one backbone whose duals thus have one boundary component. We refer to these duals as *unicellular maps*.

A bijection is derived by Chapuy in [14] between a unicellular map of genus g together with a distinguished trisection and a unicellular map of genus $g - k$, together with $2k + 1$ labeled vertices. This is facilitated by successively slicing a vertex into three new vertices via the distinguished trisection. Since the trisection can persist through finitely many slicings, say k times, the process produces $2k + 1$ labeled vertices, reducing the genus of the map by exactly k . Gluing is the inverse to slicing: gluing a set of $2k + 1$ labeled vertices in a unicellular map increases the genus of the map by k . By construction, both, slicing and gluing processes preserve unicellularity.

In a unicellular map we have two orders, $<_{\gamma}$ and $<_{\sigma}$. The former accounts for the length of the tour of the boundary component and the latter is induced by counterclockwise rotation around a vertex, i.e. from i to $\sigma(i)$. A half-edge τ is called a *trisection* if $\tau \leq \sigma^{-1}(\tau)$ and τ is not the minimum half-edge of the vertex. A trisection marks the nontrivial order discrepancies between σ and γ . It is topologically relevant, since the number of trisections in any unicellular map having genus g equals exactly $2g$ [14]. In particular, any unicellular map having genus 0, i.e. a planar tree, has no trisections.

Let \bar{v} be a vertex with trisection τ . We slice as follows [14]: let a_1 denote

the minimum half-edge of \bar{v} , let a_2 denote the minimum half-edge between a_1 and $a_3 = \tau$ (rotating clockwise), that is larger than a_3 with respect to $<_\gamma$. Iterating this means to further slice v_3 until τ becomes its minimum half-edge and we derive the mapping $\Xi(\mathbf{m}_g, \tau) = (\mathbf{m}_{g-k}, v_1, \dots, v_{2k+1})$, where k denotes the number of slicings. The new vertices v_1, \dots, v_{2k+1} are by construction ordered by their minimum half-edges with respect to $<_\gamma$. Furthermore, Ξ is bijective.

The process can be reversed, i.e., $\Lambda(\mathbf{m}_{g-k}, v_1, \dots, v_{2k+1}) = (\mathbf{m}_g, \tau)$, derived by successively gluing v_1, \dots, v_{2k+1} in the following fashion: we first glue the last three vertices v_{2k-1} , v_{2k} and v_{2k+1} obtaining a new vertex, w , with an (intermediate) trisection. Continuing this, by gluing the next two vertices with w , produces a unicellular map \mathbf{m}_g together with the trisection τ . Iterated slicings and selecting new trisections as needed, generates a unicellular map having genus 0, i.e. a planar tree.

Definition 1. (blueprint) Suppose \mathbf{m}_g is a unicellular map of genus g having m edges together with the trisection τ . A (\mathbf{m}_g, τ) -blueprint is a sequence

$$((\mathbf{m}_g, \tau), (\mathbf{m}_{g_1}, \tau_1), \dots, (\mathbf{m}_{g_{r-1}}, \tau_{r-1}), (\mathbf{m}_0, \emptyset)),$$

where $(\mathbf{m}_{g_{i+1}}, V_{i+1}) = \Xi(\mathbf{m}_{g_i}, \tau_i)$, for $0 \leq i < r$ and $g_r = 0, \tau_r = \emptyset$.

Remark 1. In the following we shall refer to the blueprint of the dual of a pk-structure simply as the blueprint of the pk-structure. In Fig. 6 we display the blueprints of a matching having genus 2.

Remark 2. Note that for two unicellular maps having genus g , the number of blueprints is not necessarily equal. Furthermore, all blueprints can be constructed by considering all trisections at each step, respectively. The cost of finding all blueprints is exponential in g , however, the genus of RNA structures found in databases is typically less than 3.

3. λ -structures

We establish now a bijection between unicellular maps \mathbf{m}_g together with blueprints and a particular class of labeled, planar trees. The labels will allow to recover the blueprints. Note that, since some vertices may have been involved in subsequent slicing events, they are not present in the tree.

Definition 2. (λ -tree) A λ -tree, $\mathbf{m}_0^{(\sigma_v)_v}$ is a rooted, planar tree with boundary component γ , in which a vertex v carries the label $\sigma_v \in \mathbb{F}_2^r$ such that

- $\sum_{v,h} \sigma_v|_h = 2g + r$, where $1 \leq r \leq g$,
- $|\{\sigma_v \mid \sigma_v|_h = 1\}| \equiv 1 \pmod{2}$,
- $|\{v_{\sigma_v} \mid \sigma_v|_h = \sigma_v|_i = 1, h < i\}| \leq 1$ and any vertex, whose label σ_v satisfies $\sigma_v|_h = \sigma_v|_i = 1$ for $h < i$, is minimal in γ within the set of all the vertices, having label σ_v where $\sigma_v|_i = 1$.

We shall call the minimum vertex v with the property $\sigma_v|_i = 1$ transitional of level i and regular, otherwise. Note that a vertex can be transitional for multiple indices.

Lemma 1. Any unicellular map, \mathbf{m}_g , together with a blueprint, \mathbf{p} , induces a unique λ -tree, $\mathbf{m}_0^{(\sigma_v)_v}$.

Proof. Suppose we are given a unicellular map \mathbf{m}_g together with a slice-path consisting of r slicings. \mathbf{p} generates a planar tree, \mathbf{m}_0 , whose vertices we color red if they were involved in some slicing and black, otherwise. \mathbf{p} furthermore specifies the sequence of slicings and each slicing reduces topological genus by at least one, whence $1 \leq r \leq g$.

We label the black vertices by $(0, \dots, 0) \in \mathbb{F}_2^r$ and each red vertex obtains a 1 in its s th-coordinate if and only if it was involved in the s th-slicing. This produces a set of intermediate labels $(\sigma'_v)_v$.

We proceed by reducing the $(\sigma'_v)_v$: starting from $i = r$ to $i = 2$, we set

$$\sigma_v = \begin{cases} (0, \dots, x_i = 1, x_{i+1}, \dots, x_r) & \text{if } v \text{ is regular at } i \\ \sigma_v & \text{otherwise.} \end{cases}$$

Accordingly, any γ -minimal vertex of a fixed slicing retains a 1 iff it came from a previously sliced vertex and, as a result, irrespective of how many times the vertices were sliced, the total number of 1 entries in all coordinates of all labels equals:

$$\sum_{\sigma_v} |\{h \mid \sigma_v|_h = 1\}| = \sum_{i=1}^r (2\Delta g_i + 1) = 2g + r,$$

where Δg_i is the decrease of genus in the i th slicing. As a result, we have $\sum_{v,h} \sigma_v|_h = 2g + r$.

By construction, the total number of 1 entries found as sth-coordinates equals the total number of vertices obtained via the sth slicing, whence $|\{\sigma_v \mid \sigma_v|_s = 1\}| \equiv 1 \pmod{2}$.

Finally, consider two fixed indices $h < i$. Only the label of the γ -minimal vertex amongst the vertices of the i th slicing can have an entry 1 as h th coordinate, whence $|\{v \mid \sigma_v|_h = \sigma_v|_i = 1, h < i\}| \leq 1$ and the proof of the lemma is complete. In Fig. 7 we illustrate the idea of the proof. \square

Having established the mapping from unicellular maps of genus g together with blueprints into λ -trees, we proceed by showing that this mapping is bijective.

Proposition 1. *Suppose \mathbf{m}_g is a unicellular map with genus g and blueprint \mathbf{p} , then we have the bijection*

$$\xi: (\mathbf{m}_g, \mathbf{p}) \mapsto \mathbf{m}_0^{(\sigma_v)_v}.$$

Proof. Following the slice-path, \mathbf{p} , the fatgraph \mathbf{m}_g is sliced to the tree \mathbf{m}_0 with labels $(\sigma_v)_v$, $\mathbf{m}_0^{(\sigma_v)_v}$. By Lemma 1, $\mathbf{m}_0^{(\sigma_v)_v}$ is a $*$ -tree, whence $\xi(\mathbf{m}_g) = \mathbf{m}_0^{(\sigma_v)_v}$ is well-defined.

To show that ξ is bijective, let $\mathbf{m}_0^{(\sigma_v)_v}$ be an arbitrary $*$ -tree such that $\sigma_v \in \mathbb{F}_2^r$. We shall construct ξ^{-1} as follows:

Let $V_r = \{v \mid \sigma_v|_r = 1\}$ denote the set of vertices whose labels satisfy $\sigma_v|_r = 1$. Lemma 1 guarantees that the number of vertices contained in V_r is odd. Then we glue

$$\Lambda: (\mathbf{m}_0, V_r) \rightarrow (\mathbf{m}_{g_{r-1}}, \tau_{r-1}),$$

creating the trisection, τ_{r-1} , where $g_{r-1} = (|V_r| - 1)/2$.

By construction $\mathbf{m}_{g_{r-1}}$ is a unicellular map with two types of vertices: those which correspond to \mathbf{m}_0 -vertices (type 1) and a unique new vertex obtained by gluing the set V_r (type 2). We now label the $\mathbf{m}_{g_{r-1}}$ -vertices with labels σ_v^{r-1} using the labels σ_v of \mathbf{m}_0 as follows: For any type 1 vertex we set $\sigma_v^{r-1}|_i = \sigma_v|_i$ for $1 \leq i \leq r-1$ and $\sigma_v^{r-1}|_r = 0$. The unique type 2 vertex obtains the label

$$\sigma_v^{r-1}|_i = \sum_{v \in V_r} \sigma_v|_i, \quad \forall 1 \leq i \leq r-1, \quad \sigma_v^{r-1}|_r = 0.$$

In Fig. 8 we illustrate the idea of the proof. \square

A λ -tree corresponds to a unicellular map of genus g , together with a blueprint. As a planar tree it corresponds by Poincaré duality to a non-crossing matching. We call the edge-labeled Poincaré dual of a λ -tree a λ -matching, $M_0^{(\sigma_a)_a}$, whose edge-labels $(\sigma_a)_a$ are induced by the vertex labels of the λ -tree (σ_v) as follows: there are $n + 1$ vertices and n edges in a planar tree, whence a vertex corresponds uniquely to an edge (except of the root). To this end, we consider the edge which connects v to its parent, see Fig. 9. Note that, by construction, the root of the planar tree is never involved in slicings and accordingly labeled $0 \in \mathbb{F}_2^r$.

A λ -matching induces an arc-labeled secondary structure by insertion of unpaired vertices. To facilitate this we employ the index map $I: \mathbb{N}_0 \rightarrow \mathbb{N}_0$, $i \rightarrow j$. We interpret I as to specify how many unpaired vertices follow the $M_0^{(\sigma_a)_a}$ -vertex, i . Starting with the left-endpoint of the rainbow and we have $I^{-1}(j + 1) - I^{-1}(j) - 1$ such vertices.

Definition 3. A λ -structure $S^{(\sigma_a)_a}$ is a noncrossing diagram with labeled arcs, σ_a , that when removing all unpaired vertices, induces a λ -matching.

Note that slicing induces transpositions of the boundary component. Thus, taking the Poincaré dual, these transpositions permute the backbone of $M_g^{(\sigma_a)_a}$. Let $\rho = \rho(\mathbf{p})$ denote resulting permutation of the backbone, where $\rho(0) = 0$, then $I^\rho(h) = I(\rho^{-1}(h))$ is the induced index map for the pk-structure corresponding to the λ -structure and we have the following situation, see Fig. 10.

$$\begin{array}{ccccc}
 (S_{g,n}, \mathbf{p}) & \text{---} & (M_g, \mathbf{p}, I^\rho) & \xrightarrow{\text{dual}} & (\mathbf{m}_g, \mathbf{p}, I^\rho) \\
 \downarrow & & \downarrow & & \downarrow \\
 S^{(\sigma_a)_a} & \text{---} & (M_0^{(\sigma_a)_a}, I) & \xrightarrow{\text{dual}} & (\mathbf{m}_0^{(\sigma_v)_v}, I)
 \end{array}$$

4. The grammar

We can now derive the CFG for λ -structures, *RNAFeatures**. The key issue here will be the compatibility of the arc labels with the production rules.

Let us briefly review the CFG for secondary structures, implied by the recursion of [24], i.e. the tuple $G = (V, \Sigma, R, I)$, where V, Σ, S are sets of nonterminal symbols, terminal symbols and production rules, respectively.

A production rule is a mapping from V to $(\Sigma \cup V)^*$, where the asterisk represents the Kleene star operation and S is an initial symbol. A derivation starts with the initial symbol I , and in each step, replaces one of the non-terminal symbols by one of the productions. For secondary structures, the grammar consists of $V = \{S, L\}$, $\Sigma = \{d, d', s\}$, and the two production rules are

$$S \rightarrow \epsilon \quad \text{or} \quad LS \quad \text{or} \quad sS, \quad L \rightarrow dSd'. \quad (1)$$

Here $S \rightarrow \epsilon$ is the ϵ -production. The terminal symbols d, d' represent pairs of vertices forming arcs and s an unpaired vertex, respectively. The nonterminal symbol S denotes an arbitrary structure and L is an irreducible structure, i.e. covered by an external arc. Via the initial symbol S this grammar is unambiguous, i.e. each structure has a unique derivation path.

We shall now extend the grammar to λ -structures. To this end we introduce new nonterminals, $S_{(t_i)_i}^{(\ell_i)_i}$ and $L_{(t_i)_i}^{(\ell_i)_i}$. The nonterminal symbols stand for sets of substructures $N = \{S, L\}$, with labeled arcs $(\sigma_a)_a$ such that $\ell_h = \sum_{\sigma_a \in N} \sigma_a|_h$, $1 \leq h \leq r$. Aside from s we have new terminals $\{d^{\sigma_a}, d'^{\sigma_a}\}$ where $d^{\sigma_a}, d'^{\sigma_a}$ denote σ_a -labelled arcs. We set $t_i = 1$ if and only if the corresponding nonterminal contains a transitional arc for i and $t_i = 0$, otherwise.

By construction, the ℓ_h will evolve depending on the presence of transitional arcs, that is $(\ell_h)_h$ may change in more than one coordinate if a transitional arc is present, or in at most one coordinate, otherwise. $(t_i)_1^r$ can be viewed as an indicator determining which production rule applies.

Decomposition: this means to decompose a nonterminal symbol $S_{(t_i)_i}^{(\ell_i)_i}$ into one left- and one right-symbol, together with consistent labelings. A labeling for $S_{(t_i)_i}^{(\ell_i)_i} \rightarrow sS_{(p_i)_i}^{(\rho_i)_i}$ is consistent if $\ell_i = \rho_i$ and $t_i = p_i$. Furthermore, consistency for $S_{(t_i)_i}^{(\ell_i)_i} \rightarrow L_{(p_i)_i}^{(\rho_i)_i} S_{(q_i)_i}^{(\xi_i)_i}$ means

- (a) $t_i = 1$ and $\rho_i > 0$ implies $p_i = 1$,
- (b) $t_i = 1$ and $\rho_i = 0$ implies $\xi_i > 0$ and $q_i = 1$
- (c) $t_i = 0$ implies $p_i = q_i = t_i = 0$.

This is a consequence of transitional arcs being induced by transitional vertices by means of Poincaré duality and the latter being minimal in the boundary component. As a result a transitional arc needs to be placed in the left-nonterminal as long as the latter contains any arcs a with the property $\sigma_a|_i = 1$.

As a result we have the production rule:

$$S_{(t_i)_i}^{(\ell_i)_i} \rightarrow \epsilon \quad \text{or} \quad L_{(p_i)_i}^{(\rho_i)_i} S_{(q_i)_i}^{(\xi_i)_i} \quad \text{or} \quad sS_{(t_i)_i}^{(\ell_i)_i}, \quad (2)$$

satisfying $\rho_i + \xi_i = \ell_i$, for $\rho_i, \xi_i \geq 0, 1 \leq i \leq r$. In case of $t_i = 1$, if $\rho_i > 0$, we have $p_i = 1$ and $q_i = 0$. Otherwise, if $\rho_i = 0$, the i th trans-arc is distributed to the right and we have $p_i = 0$ and $q_i = 1$. Finally, if $t_i = 0$, both, $p_i = q_i = 0$, see Fig. 11.

Arc-removal: this means to act on the nonterminal $L_{(t_i)_i}^{(\ell_i)_i}$ by removing its outer arc, a . The key point here is to provide this arc with a label, σ_a consistent with the label associated with the generated nonterminals. Given $L_{(t_i)_i}^{(\ell_i)_i}$ any arc-removal based production rule generates the nonterminal $S_{(t'_i)_i}^{(\ell'_i)_i}$, where $t'_i \leq t_i$ and $\ell_i - \ell'_i \leq 1$. This implies for the label of the removed arc $\sigma_a|_h = (\ell_h - \ell'_h)$, $1 \leq h \leq r$. In case a is a trans-arc for s_1, \dots, s_k , we have $0 = t'_{s_j} < t_{s_j} = 1$, or $0 = t'_j = t_j$, otherwise. Accordingly, we have:

$$L_{(t_i)_i}^{(\ell_i)_i} \rightarrow d^{\sigma_a} S_{(p_i)_i}^{(\ell_i)_i - \sigma_a} d'^{\sigma_a}. \quad (3)$$

The terminal symbols $\Sigma = \{d^{\sigma_j}, d'^{\sigma_j}, s\}$, nonterminal symbols $V = \{S_{(t_i)_i}^{(\ell_i)_i}, L_{(t_i)_i}^{(\ell_i)_i}\}$, production rules R given by eq. [2], eq. [3], together with the initial condition $I = S_{(t_i)_i}^{(\ell_i)_i}$ with $\sum_i \ell_i = 2g + r$ and $t_i = 1$, for $1 \leq i \leq r$, define the context-free grammar $G = (V, \Sigma, R, I)$.

Proposition 2. *The context-free grammar G discussed above generates all λ -structures uniquely, or equivalently, all pk -structures together with a blueprint.*

Proof: Consider a nonterminal symbol $S_{(t_i)_i}^{(\ell_i)_i}$ where the ℓ_i are positive odd integers, with $\sum_i \ell_i = 2g + r$ and $t_i = 1$. Any nonterminal symbol is uniquely generated by G since the product of each production rule is unique. Furthermore, no labeled arc with $\sigma_a|_h = 1$ can be further added to the structure when t_h is set to be 1. This guarantees that the trans-arc of level h is minimal among those arcs with $\sigma_a|_h = 1$. Thus the structures generated from $S_{(t_i)_i}^{(\ell_i)_i}$ are λ -structures. As for its generative power, an arbitrary λ -structure, is translated via G into a unique sequence of terminal symbols, whence the proposition.

The grammar G for λ -structure is designed based on the simple grammar for secondary structure by assigning labels to symbols. However, the latter evaluates arcs against unpaired vertices, lacking the capability of differentiating base pairs and unpaired bases in the context of specific loop types.

We accordingly present a grammar for λ -structures, *RNAFeatures** used in practice that differentiates loop types, that is an extension of the grammar *RNAFeatures* [44]. The grammar allows us to differentiate loop-types. First

we have the terminal symbols $\Sigma = \{A, C, G, U, R\}$ representing the respective RNA nucleotides and the endpoint of the rainbow. The nonterminal symbols are

$$\begin{aligned} V = \{ & \lambda_g\text{-structure}, \textit{ExteriorLoop}_{(t_i)_i}^{(\ell_i)_i}, \textit{Stack}_{(t_i)_i}^{(\ell_i)_i}, \textit{Weak}_{(t_i)_i}^{(\ell_i)_i}, \\ & \textit{HairpinLoop}_{(0)_i}^{\sigma_a}, \textit{InteriorLoop}_{(t_i)_i}^{(\ell_i)_i}, \textit{BulgeLeft}_{(t_i)_i}^{(\ell_i)_i}, \\ & \textit{BulgeRight}_{(t_i)_i}^{(\ell_i)_i}, \textit{MultiLoop}_{(t_i)_i}^{(\ell_i)_i}, \textit{MLComponent}_{(t_i)_i}^{(\ell_i)_i}, \\ & \textit{MLComponents}_{(t_i)_i}^{(\ell_i)_i}, \textit{SingleStrand}\}. \end{aligned}$$

The nonterminal symbols are referenced via both: their semantics and their arc-label information of the class of structures it represents. λ_g -structure are λ -structures corresponding to a pk-structure of genus g , together with a blueprint. For nonterminal symbols the $N_{((t_i)_i)}^{(\ell_i)}$, where

$$N = \{\textit{ExteriorLoop}, \textit{Stack}, \dots, \textit{MLComponents}\},$$

the index (ℓ_i) represents the sum of the arc-labels contained in the class, such that $t_i = 1$ if and only if the transitional arc of level i is present. We distinguish the following classes:

- *ExteriorLoop* denotes the class that is immediately nested in the rainbow,
- *Stack* denotes the class covered by parallel arcs, while *Weak* represents the class covered by a isolated arc,
- *HairpinLoop* represents the class containing no arcs, immediately nested in *Weak*, while *InteriorLoop*, *BulgeLeft*, *BulgeRight* have one and *MultiLoop* have more than one such arc. The classes *InteriorLoop*, *BulgeLeft* and *BulgeRight* differ in the number of unpaired bases in the two intervals formed between the outer arc and the immediately nested arc. *BulgeLeft* has the interval to the right empty and the intervals to the left non-empty, *BulgeRight* is defined analogously and in *InteriorLoop* both these intervals are non-empty,
- *MLComponents* and *MLComponent* are classes contained in a multi-loop: *MLComponent* has exactly one braching and *MLComponents* has at least two,

- *SingleStrand* denotes the class consisting of a sequence of unpaired bases.

The nonterminal symbols are listed in Fig. 12

- *ini_r*: here we have $\sum_{0 < i \leq r} \ell_i + \sigma_a = 2g + r$. In case of $\sigma_a \neq (0)_i$, the rainbow is always a transitional arc, hence $t_i = 0$ if $\sigma_a|_i \neq 0$,
- *exl₁* emits the signal ϵ of the λ -structure,
- *exl₂*($x, (\ell_i)_i, (t_i)_i$) cuts off the leftmost unpaired base and keeps all the labels as well as transitional arcs in the right nonterminal, where $x \in \{A, C, G, U\}$. *co₁*($x, (\ell_i)_i, (t_i)_i$), *ss₁*(x) and *ss₂*(x) are defined analogously,
- *exl₃*((ℓ_i)_{*i*}, (ρ_i)_{*i*}, (t_i)_{*i*}) decompose into two classes, distributing any labeled arcs and in particular any trans-arcs, whence $(\ell_i)_i = (\rho_i)_i + (\xi_i)_i$. In case of $t_i = 1$, if $\rho_i > 0$, we have $p_i = 1$ and $q_i = 0$. Otherwise, if $\rho_i = 0$, the *i*th trans-arc is distributed to the RHS and we have $p_i = 0$ and $q_i = 1$. Finally, if $t_i = 0$, both, $p_i = q_i = 0$. The rules *cs₁*((ℓ_i)_{*i*}, (ρ_i)_{*i*}, (t_i)_{*i*}), *cs₂*((ℓ_i)_{*i*}, (ρ_i)_{*i*}, (t_i)_{*i*}) and *cs₃*((ℓ_i)_{*i*}, (ρ_i)_{*i*}, (t_i)_{*i*}) are defined analogously,
- *st₁*($x, y, (\ell_i)_i, (t_i)_i, \sigma_1$) remove a base pair (x, y) with the arc label σ_a , $x, y \in \{A, C, G, U\}$, inducing the labels $(\ell_i)_i - \sigma_a$ for the class produced. If $\sigma_a|_i \neq 0$ and $t_i = 1$, the arc $(x^{\sigma_a}, y^{\sigma_a})$ is a trans-arc of level *i* and we set $p_i = 0$ accordingly. The rules *st₂*($x, y, (\ell_i)_i, (t_i)_i, \sigma_a$), *hl*(x, y, σ_a), *il*($x, y, (\ell_i)_i, (t_i)_i, \sigma_a$), *bl*($x, y, (\ell_i)_i, (t_i)_i, \sigma_a$), *br*($x, y, (\ell_i)_i, (t_i)_i, \sigma_a$) as well as *ml*($x, y, (\ell_i)_i, (t_i)_i, \sigma_a$), $x, y \in \Sigma$ are defined analogously.

We display the production rules in Fig. 13.

Let S be a structure and θ be an RNA sequence. A CFG allows to assign a score to a pair (S, θ) , derived from scores assigned to the production rules encountered. Namely, for each (S, θ) the CFG provides a unique parse tree and the score of (S, θ) is then computed by means of the scores of the respective production rules. Scores based on a thermodynamic energy model produces the minimum free energy (mfe) model [45, 46, 47] and a probabilistic scoring scheme derived from data bases of RNA sequences and structures leads to a SCFG $G = (V, \Sigma, R, S, P)$. In the latter case, sequence–structure pairs of the data set allow to assign scores to the production rules, which in turn induce a score of an arbitrary pair (S, θ) .

5. The SCFG of RNA pseudoknot structures

Stochastic context-free grammars (SCFG) were developed for RNA secondary structures [48, 49]. They have been employed for the *ab initio* structure prediction as well as the analysis of observables such as base pairing probabilities or loop patterns [50, 51, 52]. A SCFG is derived from a CFG by associating to each production rule r a probability p_r . That is, for any nonterminal A , suppose that r_1, \dots, r_k are all derivations of A , then we have $\sum_i p_i = 1$.

Let S_g denotes an RNA structure of genus g and θ be a sequence. For $g = 0$, S_0 is a secondary structure and can be parsed by *RNAFeatures*. We construct the parse tree of (S_0, θ) and record the frequency of the production rules. In case of $g > 0$, for any S_g -blueprint, $\mathbf{p}(S_g)$, we consider the correspondence

$$(S_g, \mathbf{p}(S_g), \theta) \rightarrow (S_{\mathbf{p}(S_g)}^{(\sigma_a)_a}, \theta_{\mathbf{p}(S_g)}), \quad (4)$$

where $S_{\mathbf{p}(S_g)}^{(\sigma_a)_a}$ is a λ -structure and $\theta_{\mathbf{p}(S_g)}$ is the modified sequence. We construct the parse tree of *RNAFeatures** of $(S_{\mathbf{p}(S_g)}^{(\sigma_a)_a}, \theta_{\mathbf{p}(S_g)})$ and record the frequency of the encountered production rules.

If there are m S_g -blueprints, the contribution of each pair $(S_{\mathbf{p}(S_g)}^{(\sigma_a)_a}, \theta_{\mathbf{p}(S_g)})$ is normalized by $1/m$ and the frequency of a production rule R equals $f(R)/m$ where $f(R)$ is the frequency of R being applied. After processing all structures and sequences, the probability of a production rule is computed as the quotient of its frequency over the sum of all frequencies of its respective derivation.

We next consider $\mathbb{P}(S_g, \theta|M)$, the normalized score of an arbitrary structure, S_g , over an arbitrary sequence, θ , as a function of the data set M . We set $\mathbb{P}(S_g, \theta|M) = \sum_{\mathbf{p}(S_g)} \mathbb{P}(S_{\mathbf{p}(S_g)}^{(\sigma_a)_a}, \theta_{\mathbf{p}(S_g)}|M)$, where $\mathbb{P}(S_{\mathbf{p}(S_g)}^{(\sigma_a)_a}, \theta_{\mathbf{p}(S_g)}|M)$ is calculated by the multiplying the probabilities of the production rules in the parse tree of $(S_{\mathbf{p}(S_g)}^{(\sigma_a)_a}, \theta_{\mathbf{p}(S_g)})$ and the sum is taken over the set of all S_g -blueprints.

The $\mathbb{P}(S_g, \theta|M)$ induce $\mathbb{P}(S_g|M) = \sum_{\theta \in \mathcal{Q}_4^n} \mathbb{P}(S_g, \theta|M)$. Using the fact

that for fixed blueprint, \mathbf{p} , we have a bijection $\beta_{\mathbf{p}}: Q_4^n \rightarrow Q_4^n$, we derive

$$\begin{aligned}\mathbb{P}(S_g|M) &= \sum_{\theta_{\mathbf{p}(S_g)} \in Q_4^n} \sum_{\mathbf{p}(S_g)} \mathbb{P}(S_{\mathbf{p}(S_g)}^{(\sigma_a)^a}, \theta_{\mathbf{p}(S_g)} | M) \\ &= \sum_{\mathbf{p}(S_g)} \mathbb{P}(S_{\mathbf{p}(S_g)}^{(\sigma_a)^a} | M).\end{aligned}$$

Furthermore the SCFG facilitates the sampling of structures $S_g \in \mathbb{S}_{g,n}$ with Boltzmann probability, i.e., $\mathbb{P}(S_g) = \mathbb{P}(S_g|M) / \sum_{S \in \mathbb{S}_{g,n}} \mathbb{P}(S|M)$. We employ *RNAFeatures** to Boltzmann sample a λ -structure S^* , and reconstruct S_g as well as the associated blueprint, $\mathbf{p}(S_g)$: fixing S_g , suppose there are m blueprints $\mathbf{p}_i(S_g)$, these correspond to the m distinct λ -structures S_i^* . Let \mathbb{S}_n^* be the collection of all such S_i^* . In view of $S_i^* \rightarrow (S_g, \mathbf{p}_i(S_g))$ being a bijection, we have

$$\sum_{S_g \in \mathbb{S}_{g,n}} \mathbb{P}(S_g|M) = \sum_{S_g \in \mathbb{S}_{g,n}} \sum_{\mathbf{p}(S_g)} \mathbb{P}(S_{\mathbf{p}(S_g)}^{(\sigma_a)^a} | M) = \sum_{S^* \in \mathbb{S}_n^*} \mathbb{P}(S^* | M)$$

Therefore, in view of $\mathbb{P}(S_g) = \sum_{\mathbf{p}_i(S_g)} \mathbb{P}(S_i^*)$ and

$$\sum_{\mathbf{p}_i(S_g)} \mathbb{P}(S_i^*) = \frac{\sum_{\mathbf{p}_i(S_g)} \mathbb{P}(S_i^* | M)}{\sum_{S^* \in \mathbb{S}_n^*} \mathbb{P}(S^* | M)} = \frac{\mathbb{P}(S_g | M)}{\sum_{S_g \in \mathbb{S}_{g,n}} \mathbb{P}(S_g | M)}$$

we have $\mathbb{P}(S_g) = \sum_{\mathbf{p}_i(S_g)} \mathbb{P}(S_i^* | M) / \sum_{S^* \in \mathbb{S}_n^*} \mathbb{P}(S^* | M)$.

As for runtime complexity, the Boltzmann sampler of secondary structures using such a SCFG has runtime of $O(n^2)$, if structures are ranked according to the sequential order. This can be improved to $O(n \log n)$ by ranking the structures with the Boustrophedon order [53, 9]. For fixed topological genus, the label enhancement of *RNAFeatures** by construction increases only the number of derivations. Thus the Boltzmann sampler for λ -structures has runtime complexity $O(n \log n)$.

6. Discussion

RNA secondary structures have been widely studied using the classic CFG implied by the recursions in [16, 24]. Using RNA sequence and structure data sets, this CFG induces a SFCG. This SCFG has been applied for the *ab initio* prediction of RNA secondary structures [28, 50], calculating the likelihood

of a sequence generated by the model SCFG [51], and detecting RNA genes using comparative sequence analysis [52]. Furthermore, the generation of random secondary structures according to the native distributions of certain families of RNA structures is studied in [9] in order to identify key structural motifs in biological RNA secondary structures. This paper allows to extend these studies to pk-structures. To the best of our knowledge *RNAFeatures** is the first unambiguous context-free grammar for pk-structures. Related work has been done by Freiermuth and Nebel on a CFG of certain, colored trees. These are trees that can be mapped to unicellular maps, however, not bijectively.

A particular, multiple context-free grammar (MCFG) for pk-structures is introduced in [11]. This grammar employs a vector of nonterminal symbols referencing a substructure with a gap. The time and space complexity for Boltzmann sampling pk-structures with this grammar are $O(n^4)$ and $O(n^2)$, respectively. The grammar furthermore implies a folding algorithm having relatively high computational cost [11]: $O(n^6)$ time and $O(n^4)$ space complexity, using the thermodynamic model. Restricting to certain types of pk-structures the time complexity can be reduced to $O(n^4)$ time complexity [54]. In any case, specifying the output space of the MCFG, i.e. what types of structures the grammar actually generates is subtle, see [17] for details.

The passage from MCFG to CFG has implications for time and space complexity: an MCFG contains rules of the type $A \rightarrow B_1 C_1 B_2 C_2$ where $\langle B_1, B_2 \rangle$, $\langle C_1, C_2 \rangle$ are vectors of nonterminals. It requires a minimum of three loop-indices to execute these rules in the sampling or folding process. In contrast, the production rule in any CFG being in Chomsky normal form can be written as $A \rightarrow BC$, which requires only one loop index, reducing the time complexity by $O(n^2)$. As for space complexity this passage obsoletes using vectors, reducing the space complexity for $O(n)$ and $O(n^2)$ for sampling and folding, respectively.

The key idea facilitating this passage lies in mapping the information of crossing arcs into the labels of noncrossing arcs, i.e. passing from pk-structures to λ -structures. Accordingly, the bijection of eq. 4 is the crucial point. The fact that the bijection is genuinely topology-based allows *RNAFeatures** to extend the classic CFG of secondary structures, generating any pk-structure.

In the following we present the expectation and variance of some important parameters related to the structural features of pk-structures generated by the induced SCFG of *RNAFeatures**. These parameters are also used to

analyze Boltzmann sampled secondary structures in comparison with native structures in [9]. Our SCFG is trained via 90 single stranded tRNA sequence-structure pairs, taken from the Nucleic Acid Database (NDB) [55, 56], considering only Watson-Crick and wobble base pairs. In this training set, we find 16 structures having genus 0, 73 of genus 1 and 1 structure of genus 2. The average sequence length is 75.21 nucleotides. We analyze the following five random variables: bp , the number of base pairs, st_n , the number of stacks, st_ℓ , the average length of a stack, hp_n , the number of hairpin and finally hp_ℓ , the average length of a hairpin loop. Here a stack of length k is the maximal sequence of parallel arcs $\{(i, j), (i + 1, j - 1), \dots, (i + k - 1, j - k + 1)\}$. A hairpin loop of length $j - i - 1$ is a substructure, which consists of a base pair (i, j) and the unpaired bases $i + 1, \dots, j - 1$.

The statistics is obtained by Boltzmann sampling of 10^5 structures over 76 nucleotides via the SCFG. For reference purposes we furthermore sample 10^5 random structures over 76 nucleotides, having fixed genus 0 and 1, respectively, with uniform probability [57]. We summarize in Table 1 the respective means and variances of the above five random variables.

The sampling generates 264, 9610, 126 structures having genus 0, 1 and 2, reflecting the large quantity of genus 1 structures in the training set. By construction our SCFG does not recognize contributions of unpaired bases in hairpin loops, bulges and interior loops, which explains the systematic deviation in the lengths of hairpin-loops. Table 1 shows that the Boltzmann sampled structures are not random structures of either genus zero or one.

The implication of the CFG to the folding of pk-structures is work in progress. Though we can fold a λ -structure over a given sequence θ in $O(n^3)$ time, the bijection producing the corresponding pseudoknotted structure does systematically alter the underlying sequence. In particular, in case of pk-structures of genus one, this consists in the transposition of two subsequent intervals and is thus manageable.

7. Acknowledgments.

We gratefully acknowledge the help of Kevin Shinpaugh and the computational support team at VBI. Many thanks to Christopher L. Barrett and Henning Mortveit for discussions.

References

- [1] O. C. Uhlenbeck, A coat for all sequences, *Nature Struct. Biol.* 5 (1998) 174–176.
- [2] P. R. Stein, M. S. Waterman, On some new sequences generalizing the catalan and motzkin numbers, *Discrete Math.* 26 (1978) 261–272.
- [3] M. S. Waterman, T. F. Smith, Rapid dynamic programming methods for RNA secondary structure, *Adv. Appl. Math.* 7 (1986) 455–464.
- [4] M. Zuker, P. Stiegler, Optimal computer folding of larger RNA sequences using thermodynamics and auxiliary information, *Nucleic Acids Res.* 9 (1981) 133–148.
- [5] I. L. Hofacker, W. Fontana, P. F. Stadler, L. S. Bonhoeffer, M. Tacker, P. Schuster, Fast folding and comparison of RNA secondary structures, *Monatsh. Chem.* 125 (1994) 167–188.
- [6] J. S. McCaskill, The equilibrium partition function and base pair binding probabilities for RNA secondary structure, *Biopolymers* 29 (1990) 1105–1119.
- [7] P. Duchon, P. Flajolet, G. Louchard, G. Schaeffer, Boltzmann samplers for the random generation of combinatorial structures, *Combinatorics, Probability and Computing* 13 (2004) 2004.
- [8] M. A. Harrison, *Introduction to Formal Language Theory*, Addison-Wesley, 1978.
- [9] M. E. Nebel, A. Scheid, F. Weinberg, Random generation of RNA secondary structures according to native distributions, *Algorithms Mol. Biol.* 6 (24) (2011) doi:10.1186/1748-7188-6-24.
- [10] W. Fontana, D. A. M. Konings, P. F. Stadler, P. Schuster, Statistics of RNA secondary structures, *Biopolymers* 33 (1993) 1389–1404.
- [11] E. Rivas, S. R. Eddy, A dynamic programming algorithm for RNA structure prediction including pseudoknots, *J. Mol. Biol.* 285 (1999) 2053–2068.

- [12] N. B. Leontis, E. Westhof, Geometric nomenclature and classification of RNA base pairs, *RNA* 7 (2001) 499–512.
- [13] R. C. Penner, M. Knudsen, C. Wiuf, J. E. Andersen, Fatgraph models of proteins, *Comm. Pure Appl. Math.* 63 (2010) 1249–1297.
- [14] G. Chapuy, A new combinatorial identity for unicellular maps, via a direct bijective approach, *Adv. Appl. Math.* 47(4) (2011) 874–893.
- [15] W. S. Massey, *Algebraic Topology: An Introduction*, Springer-Verlag, New York, 1967.
- [16] T. Smith, M. Waterman, RNA secondary structure, *Math. Biol.* 42 (1978) 31–49.
- [17] E. Rivas, S. R. Eddy, The language of RNA: A formal grammar that includes pseudoknots, *Bioinformatics* 16 (2000) 334–340.
- [18] L. Pachter, B. Sturmfels, *Algebraic Statistics for Computational Biology*, Cambridge University Press, 2005.
- [19] R. Durbin, S. R. Eddy, A. Krogh, G. Mitchison, *Biological Sequence Analysis*. Cambridge, Cambridge University Press, 1998.
- [20] K. Munch, A. Krogh, Automatic generation of gene finders for eukaryotic species, *BMC Bioinformatics* 7 (2006) 263.
- [21] L. Pachter, M. Alexandersson, S. Cawley, Applications of generalized pair hidden markov models to alignment and gene finding problems, *J. Comput. Biol.* 9 (2) (2002) 389–99.
- [22] C. Lottaz, C. Iseli, C. V. Jongeneel, P. Bucher, Modeling sequencing errors by combining hidden markov models, *Bioinformatics* 19 Suppl 2 (2003) ii103–12.
- [23] B. J. Yoon, P. P. Vaidyanathan, Structural alignment of RNAs using profile-csHMMs and its application to RNA homology search: Overview and new results., *IEEE Trans. Automat. Contr.* (Joint Special Issue on Systems Biology with IEEE Transactions on Circuits and Systems: Part-I) 53 (2008) 10–25.

- [24] M. S. Waterman, Secondary structure of single-stranded nucleic acids, *Adv. Math. (Suppl. Studies)* 1 (1978) 167–212.
- [25] R. Nussinov, G. Piecznik, J. R. Griggs, D. J. Kleitman, Algorithms for loop matching, *SIAM J. Appl. Math.* 35 (1) (1978) 68–82.
- [26] D. Kleitman, Proportions of irreducible diagrams, *Studies in Appl. Math.* 49 (1970) 297–299.
- [27] D. Sankoff, G. Leduc, N. Antoine, B. Paquin, B. F. Lang, C. R., Gene order comparisons for phylogenetic inference: Evolution of the mitochondrial genome, *Proc. Natl. Acad. Sci. USA* 89 (1992) 6575–6579.
- [28] B. Knudsen, J. J. Hein, Using stochastic context-free grammars and molecular evolution to predict RNA secondary structure, *Bioinformatics* 15 (1999) 446–454.
- [29] E. Westhof, L. Jaeger, RNA pseudoknots, *Curr. Opin. Struct. Biol.* 2 (1992) 327–333.
- [30] J. L. Chen, C. W. Greider, Functional analysis of the pseudoknot structure in human telomerase RNA, *Proc. Natl. Acad. Sci. USA* 102 (23) (2005) 8080–5.
- [31] A. Loria, T. Pan, Domain structure of the ribozyme from eubacterial ribonuclease, *RNA* 2 (1996) 551–563.
- [32] D. W. Staple, S. E. Butcher, Pseudoknots: RNA structures with diverse functions, *PLoS Biol.* 3 (2005) e213.
- [33] D. Konings, R. Gutell, A comparison of thermodynamic foldings with comparatively derived structures of 16s and 16s-like rRNAs, *RNA* 1 (1995) 559–574.
- [34] C. Tuerk, S. MacDougall, L. Gold, RNA pseudoknots that inhibit human immunodeficiency virus type 1 reverse transcriptase, *Proc. Natl. Acad. Sci. USA* 89(15) (1992) 6988–6992.
- [35] R. C. Penner, M. S. Waterman, Spaces of RNA secondary structures, *Adv. Math.* 101 (1993) 31–49.

- [36] R. C. Penner, Cell decomposition and compactification of Riemann's moduli space in decorated Teichmüller theory, in: N. Tongring, R. C. Penner (Eds.), Woods Hole Mathematics-perspectives in math and physics, World Scientific, Singapore, 2004, pp. 263–301, arXiv: math.GT/0306190.
- [37] H. Orland, A. Zee, RNA folding and large n matrix theory, Nuclear Physics B 620 (2002) 456–476.
- [38] M. Bon, G. Vernizzi, H. Orland, A. Zee, Topological classification of RNA structures, J. Mol. Biol. 379 (2008) 900–911.
- [39] L. Euler, Elementa doctrinae solidorum, Novi Comm. Acad. Sci. Imp. Petropol 4 (1752) 109–140.
- [40] J. E. Andersen, R. C. Penner, C. M. Reidys, M. S. Waterman, Topological classification and enumeration of RNA structures by genus, J. Math. Biol. 67 (5) (2013) 1261–1278.
- [41] C. M. Reidys, F. Huang, J. E. Andersen, R. C. Penner, P. F. Stadler, M. E. Nebel, Topology and prediction of RNA pseudoknots, Bioinformatics 27 (2011) 1076–1085.
- [42] M. Loebl, I. Moffatt, The chromatic polynomial of fatgraphs and its categorification, Adv. Math. 217 (2008) 1558–1587.
- [43] A. Hatcher, Algebraic Topology, Cambridge University Press, 2002.
- [44] R. Giegerich, Introduction to stochastic context free grammars, in: J. Gorodkin, W. L. Ruzzo (Eds.), RNA Sequence, Structure, and Function: Computational and Bioinformatic Methods, 1st Edition, Vol. 1097, Humana Press, 2014, Ch. 5, pp. 85–126.
- [45] D. Mathews, J. Sabina, M. Zuker, D. Turner, Expanded sequence dependence of thermodynamic parameters improves prediction of RNA secondary structure, J. Mol. Biol. 288 (1999) 911–940.
- [46] D. Mathews, M. Disney, J. Childs, S. Schroeder, M. Zuker, D. Turner, Incorporating chemical modification constraints into a dynamic programming algorithm for prediction of RNA secondary structure, Proc. Natl. Acad. Sci. USA 101 (2004) 7287–7292.

- [47] D. Turner, D. H. Mathews, NNDB: the nearest neighbor parameter database for predicting stability of nucleic acid secondary structure, *Nucl. Acids Res.* 38 (2010) D280–D282.
- [48] S. R. Eddy, R. Durbin, RNA sequence analysis using covariance models, *Nucl. Acids Res.* 22 (1994) 2079–2088.
- [49] Y. Sakakibara, M. Brown, R. Hughey, I. S. Mian, K. Sjölander, D. Underwood, R. C. and Haussler, Stochastic context-free grammars for tRNA modeling, *Nucl. Acids Res.* 22 (1994) 5112–5120.
- [50] B. Knudsen, J. J. Hein, Pfold: RNA secondary structure prediction using stochastic context-free grammars, *Nucl. Acids Res.* 31 (2003) 3423–3428.
- [51] E. Rivas, S. R. Eddy, Secondary structure alone is generally not statistically significant for the detection of noncoding RNAs, *Bioinformatics* 16 (7) (2000) 583–605.
- [52] E. Rivas, S. R. Eddy, Noncoding RNA gene detection using comparative sequence analysis, *BMC Bioinformatics* 2 (8).
- [53] Y. Ponty, Efficient sampling of RNA secondary structures from the Boltzmann ensemble of low-energy: the boustrophedon method, *J. Math. Biol.* 56 (1-2) (2008) 107–27.
- [54] R. M. Dirks, N. A. Pierce, A partition function algorithm for nucleic acid secondary structure including pseudoknots, *J. Comput. Chem.* 24 (2003) 1664–1677.
- [55] H. M. Berman, W. K. Olson, D. L. Beveridge, J. Westbrook, A. Gelbin, T. Demeny, A. R. Hsieh, S. H. amd Srinivasan, B. Schneider, The nucleic acid database: A comprehensive relational database of three-dimensional structures of nucleic acids, *Biophys. J.* 63 (751-759).
- [56] B. C. Narayanan, J. Westbrook, S. Ghosh, A. I. Petrov, B. Sweeney, C. L. Zirbel, N. B. Leontis, H. M. Berman, The nucleic acid database: new features and capabilities, *Nucleic Acids Res.* 42 (2013) D114–22.
- [57] F. Huang, M. E. Nebel, C. M. Reidys, Generation of RNA pseudoknot structures with topological genus filtration, *Math Biosci.* 245 (2) (2013) 216–225.

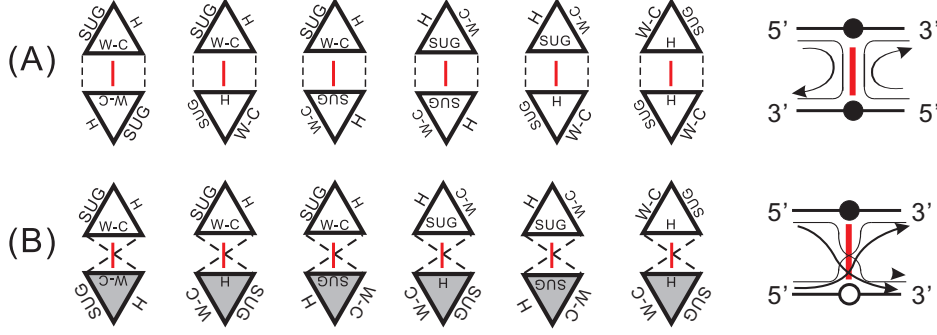


Figure 1: Natural classification of RNA base pairs [12] by faces of simplices. (A) Cis-base pairs formed by parallel triangles inducing untwisted ribbons. (B) Trans-base pairs formed by anti-parallel triangles inducing twisted ribbons.

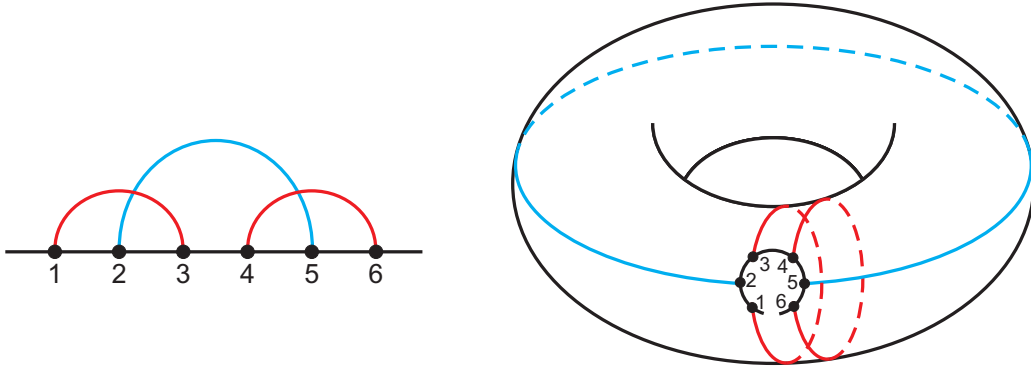


Figure 2: Crossings depend on the surface the structure is drawn on: here the diagram exhibits crossing arcs when drawn in the plane, on the torus however, the diagram can be drawn cross-free.

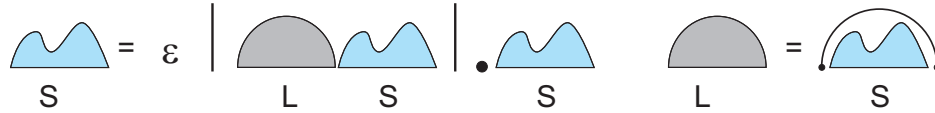


Figure 3: Illustration of the two production rules of the CFG of RNA secondary structures: decomposition (left) and arc removal (right).

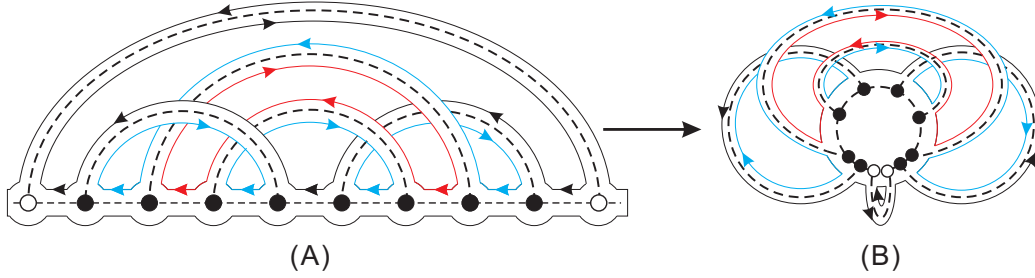


Figure 4: (A): Inflation of edges and vertices into ribbons and disks. Here we have four boundary components traversed in counter-clockwise orientation traversing the sides of any ribbon in opposite directions. (B): Collapsing the backbone of (A) into a single disc does not affect genus. Here we have $g = 1$.

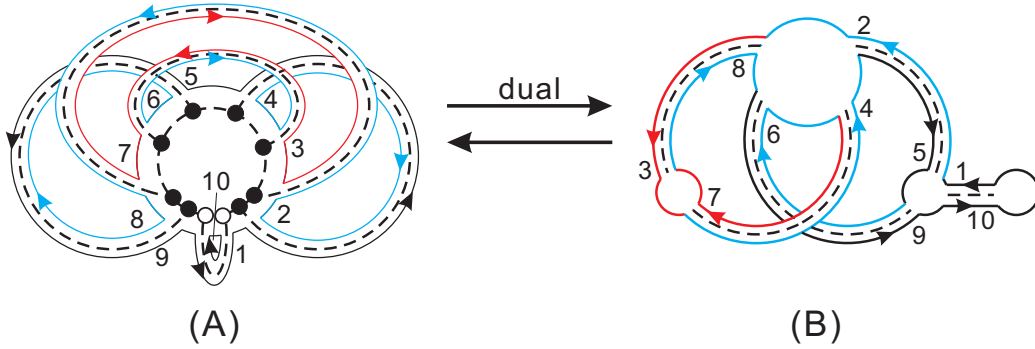


Figure 5: (A): a fatgraph represented by permutations. We have $\sigma = (1, 2, 3, 4, 5, 6, 7, 8, 9, 10)$, the unique vertex obtained by collapsing the backbone, $\alpha = (1, 10)(2, 5)(3, 8)(4, 7)(6, 9)$ the fixed-point free involution representing the ribbons, and $\gamma = (1, 5, 9)(2, 8, 6, 4)(3, 7)(10)$. (B): the Poincaré dual interchanging boundary components by vertices, hence producing a unicellular map.

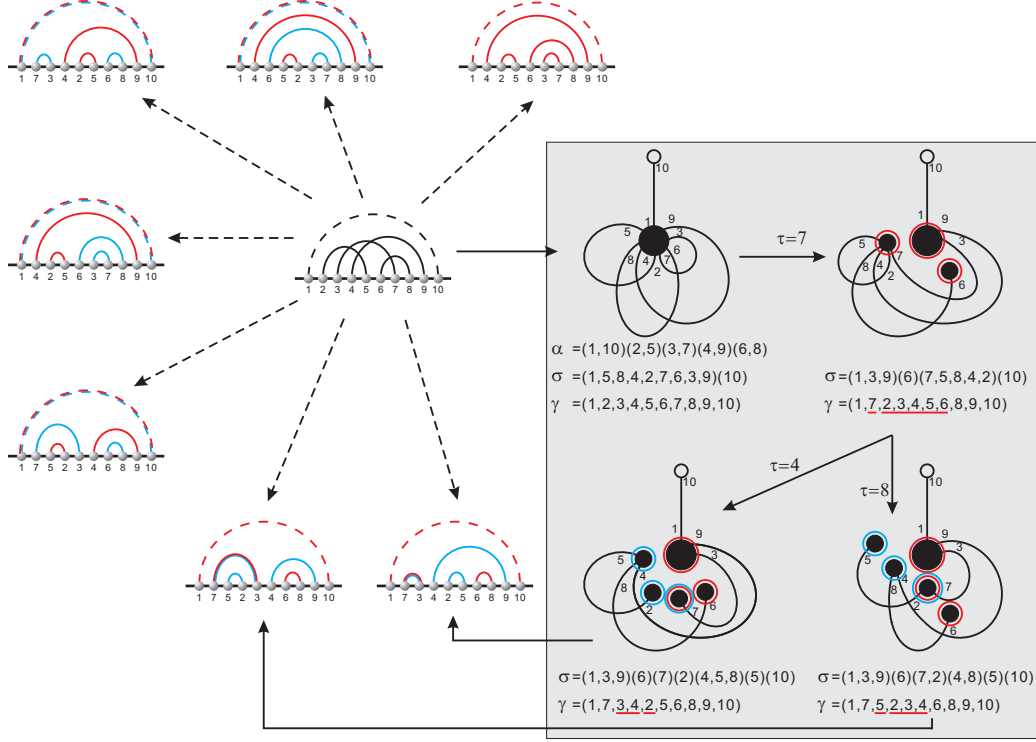


Figure 6: All blueprints of a matching of genus 2 as well as all of its associated λ -matchings induced by all possible blueprints. In the insert, we show how the bottom two λ -matchings are induced: red circles reference the first and blue circles the second slicing, respectively.

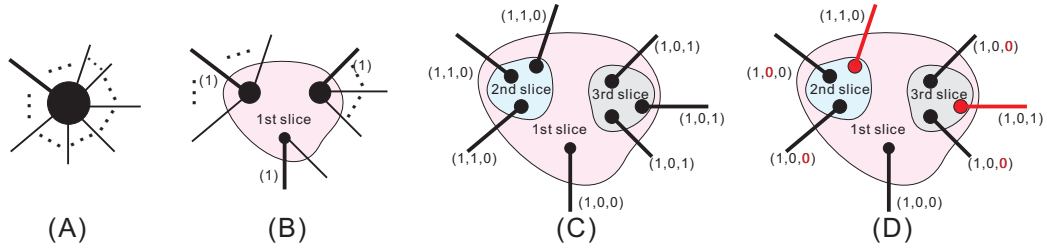


Figure 7: The idea of the proof of Lemma 1: the labels record the successive slicings and are finally reduced: only transitional vertices carry the information about previous slicing events.

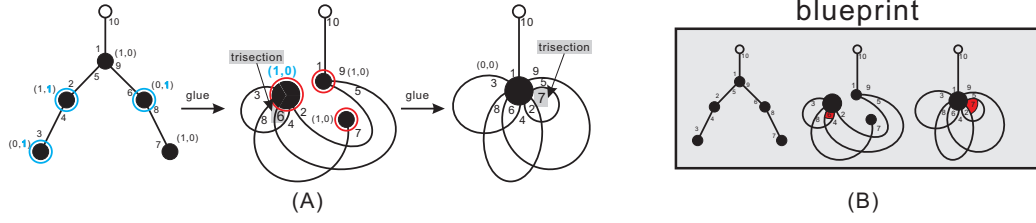


Figure 8: Reconstructing the blueprint from a λ -tree: we consider the set of all vertices such that $\sigma_v|_2 = 1$ (blue circles) and glue and relabel the vertices as in the proof of Proposition 1. This generates the unicellular map \mathbf{m}_1 with the glued vertex labeled by $(1, 0)$, carrying the distinguished trisection 8. Iteration of this process, produces \mathbf{m}_2 together with a distinguished trisection. We have thus constructed from the λ -tree the blueprint $((\mathbf{m}_2, \tau), (\mathbf{m}_1, \tau_1), (\mathbf{m}_0, \emptyset))$.

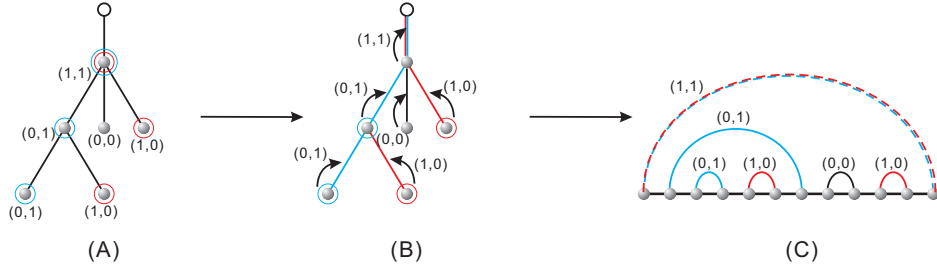


Figure 9: From a λ -tree with labeled vertices to a λ -matching with labeled arcs. (A) a λ -tree with labeled vertices. (B) shifting the labels from vertices to edges. (C) a λ -matching with labeled arcs obtained from (B) by taking the Poincaré dual.

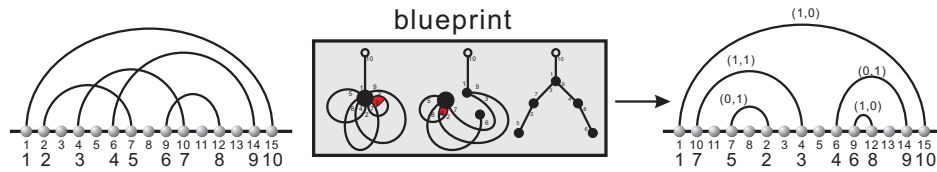


Figure 10: A pk-structure together with a blueprint maps to a unique λ -structure.

$$\begin{array}{c}
 \text{blueprint} \\
 \begin{array}{c}
 \text{Diagram of a blueprint (unicellular maps)} \\
 \rightarrow \\
 \text{Diagram of a } \lambda\text{-structure (arcs on a line)}
 \end{array}
 \end{array}$$

Figure 11: Decomposition: any transitional arc is placed in the left-nonterminal as long as the latter contains any arcs a with the property $\sigma_a|_i = 1$.

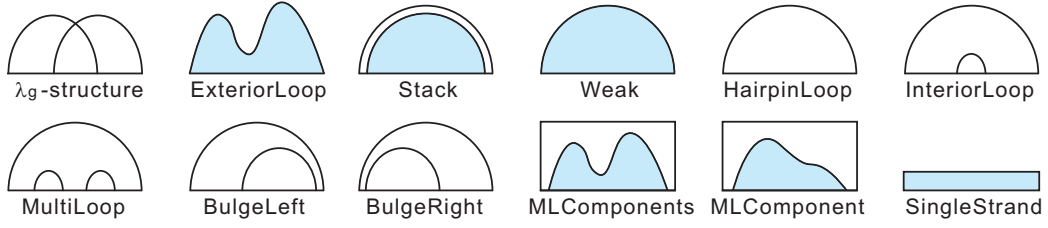


Figure 12: The nonterminal symbols of *RNAFeatures**.

Table 1: Expectation and variance of important statistic variables related to the RNA structural motifs.

bp	Input Set	Boltzmann	Uniform, $g = 1$	Uniform $g = 0$
	21.31	21.06	25.58	25.05
$E(st_n)$	5.17	5.36	22.94	22.25
$Var(st_n)$	0.56	0.91	3.98	4.25
$E(st_\ell)$	3.77	3.65	1.03	1.04
$Var(st_\ell)$	4.68	9.92	0.97	0.97
$E(hp_n)$	1.29	1.18	11.81	13.04
$Var(hp_n)$	0.50	0.21	4.27	4.29
$E(hp_\ell)$	7.16	5.55	0.47	0.51
$Var(hp_\ell)$	6.85	13.38	0.70	0.74

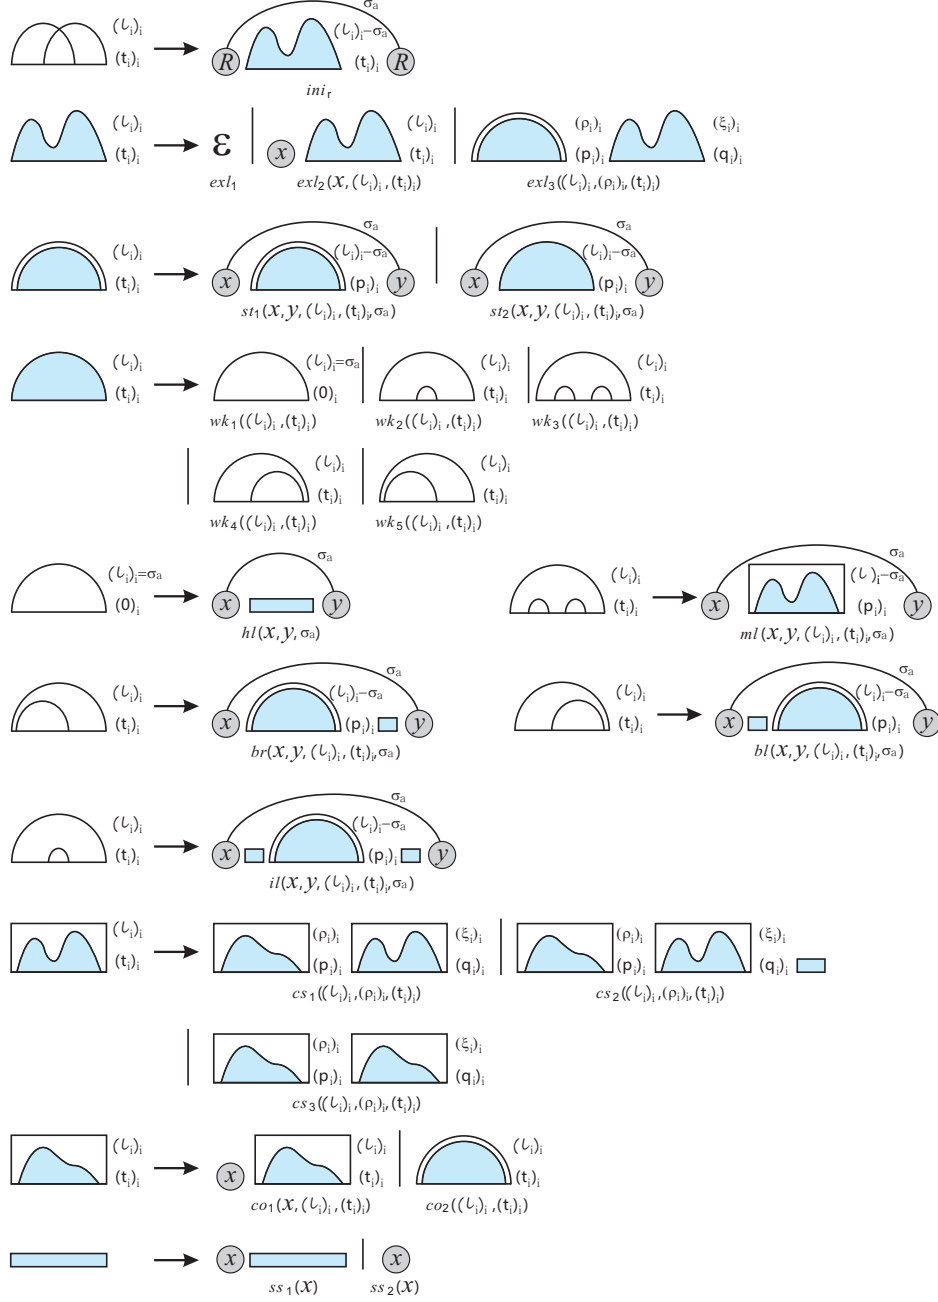


Figure 13: The production rules of $RNAFeatures^*$.